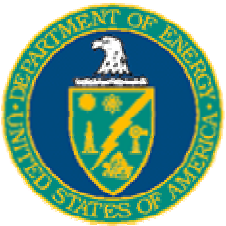


MICS Program

*What are network
needs?*

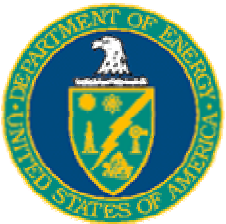
Al Geist ORNL
ESnet Steering Committee Mtg,
March 18, 2003



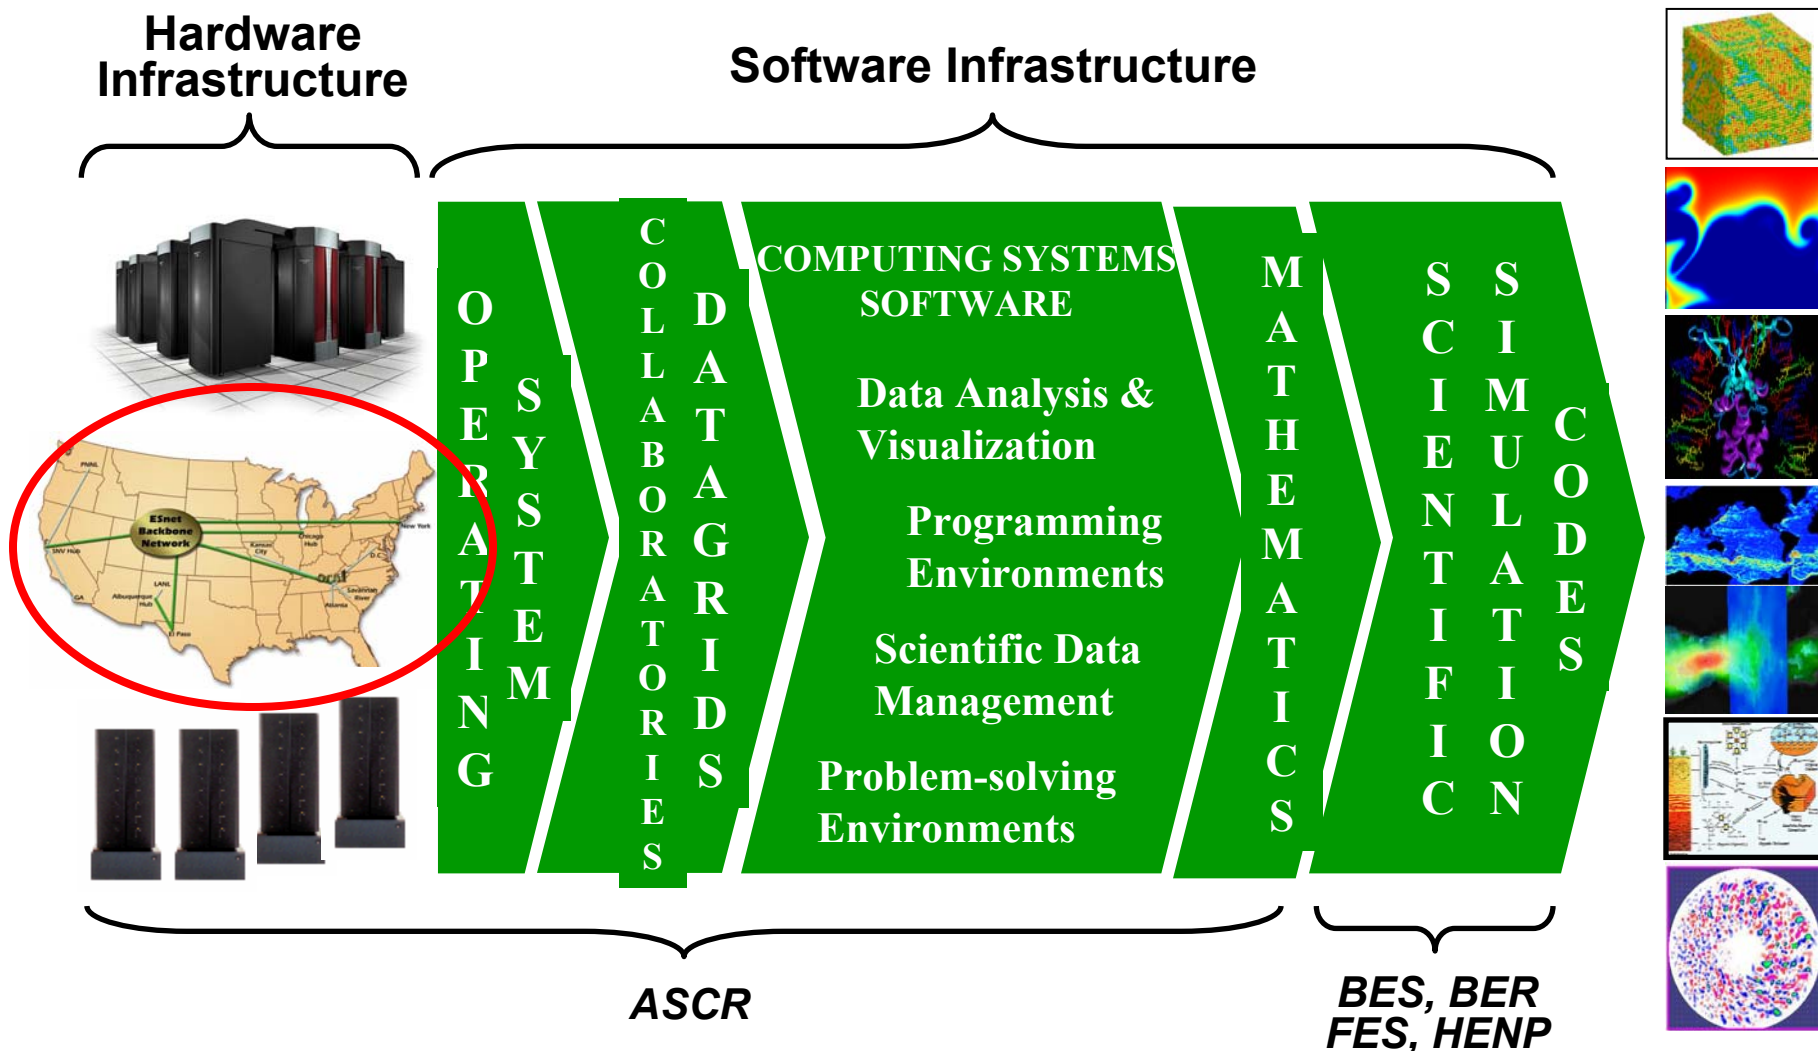
Roadmap FY02-FY08

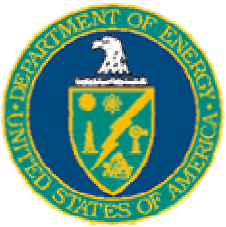
Science Drivers for Computational and Networking Needs

- **FY02 SciDAC (joint MICS and BES, BER, FES, HENP)**
- **FY03 GTL (joint MICS and BER)**
- **FY04: Next Generation Computer Architecture**
- **FY05: Ultrascale Scientific Computing Capability**
- **FY06-08: Facilities – SNS, GTL, Nanotechnology, Petascale computing**



SciDAC provides the foundation

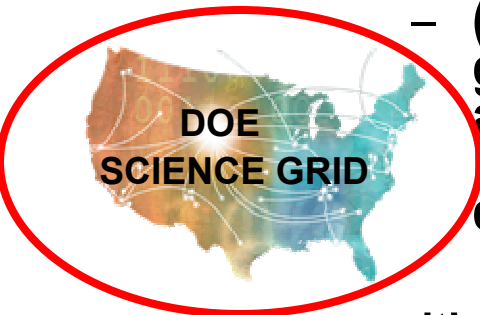




SciDAC Goals

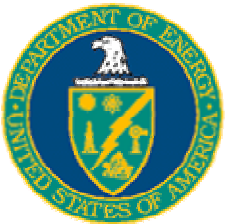
SciDAC is a pilot program for a “new way of doing science”
an INTEGRATED program to:

- (1) create a new generation of scientific simulation codes that take full advantage of the extraordinary capabilities of terascale computers
- (2) create the mathematical and computing systems software to enable scientific simulation codes to effectively and efficiently use terascale computers
- (3) create a collaboratory software environment to enable geographically distributed scientists to work effectively together as a TEAM and to facilitate remote access, through appropriate hardware and middleware infrastructure, to both facilities and data



**Cyber Security
vs Grids**

with the ultimate goal of advancing fundamental research in science
central to the DOE mission



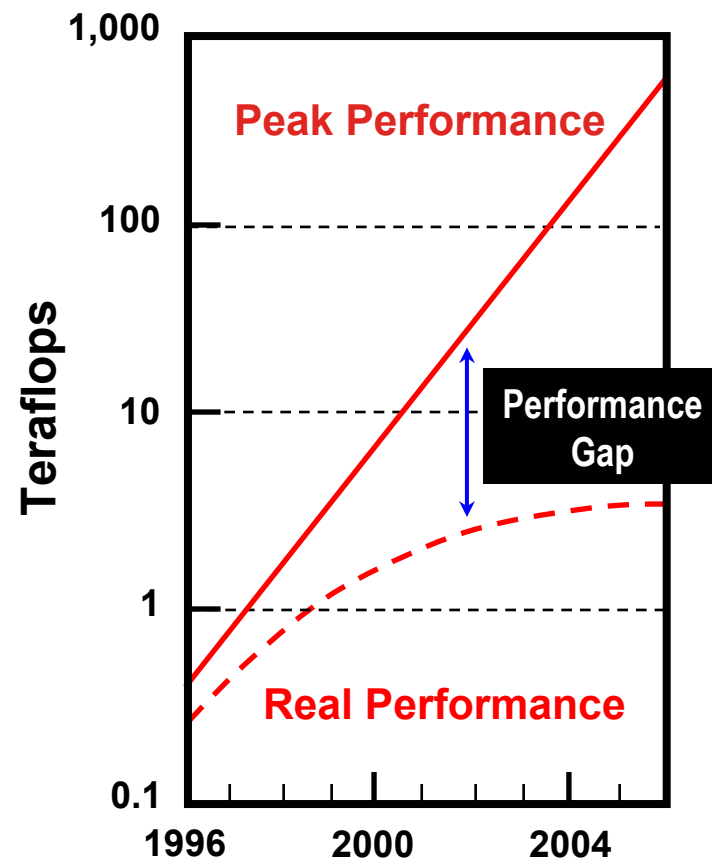
The Performance Gap

Peak performance is skyrocketing

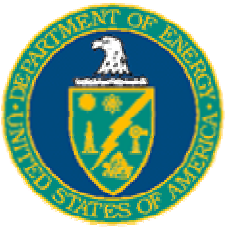
- But Efficiency for many science applications declined from 40-50% on the vector supercomputers of 1990s to as little as 5-10% on parallel supercomputers of today

Need research on ...

- Mathematical methods and algorithms that achieve high performance on a single processor and scale to thousands of processors
- More efficient programming models for massively parallel supercomputers



We need to identify a similar Gap for Scientific Networking



More Gaps to Fill

Additional computing and network resources

- initial SciDAC focus is on software, but new hardware will be needed within the next two years
- both capability and capacity computing needs are evolving rapidly

Limited architectural options available in the U.S. today

- topical computing may be a cost-effective way of providing extra computing resources
- math and CS research will play a key role

Expansion of SciDAC program

- many important SC research areas (e.g., materials/nanoscience, functional genomics/proteomics) are not yet included in SciDAC; NSRCs, GTL

DOE Genomes To Life Program

<http://www.genomes-to-life.org>

DOE awards \$103M to five large projects in GTL

Genomes to life center for molecular and cellular systems
led by ORNL and PNNL, w/ ANL, SNL,
UNC, UU

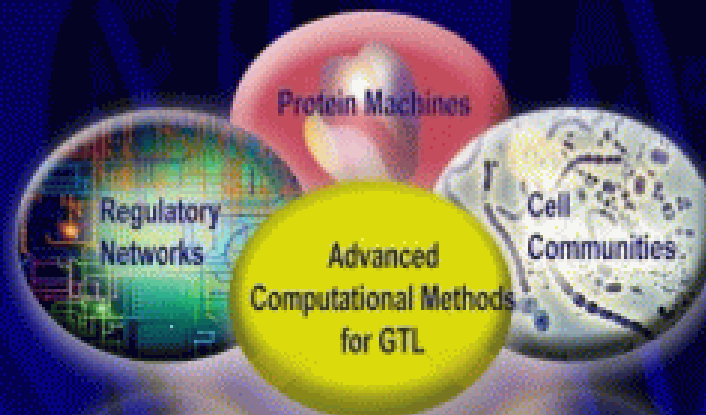
Molecular Machines to Hierarchical modeling
led by SNL and ORNL, w/ LBL, LANL,
NCGR, UCSD, UTK, UM, MSI, UCSB, UIUC

Rapid Deduction of Stress Pathways in Bacteria
led by LBL, w/ ORNL, SNL, UCB, UMC, UW,
Diversa Corporation

Microbial Ecology, Proteogenomics & Computational Optima
led by Harvard Medical School, w/ MIT,
Brigham & Women's Hospital, Massachusetts
General Hospital

Analysis of the Gene Expression of Microbial Communities
led by U. Mass., with ANL; UT, Memphis;
The Institute for Genomic Research

Understanding the Essential
Processes of Living Systems



• IMPACT •

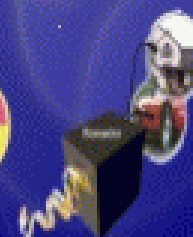
Energy
Security



Human Health
Protection

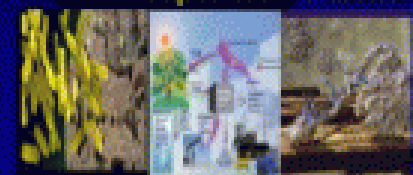


Energy
Cleanup



Tackling Problems in:

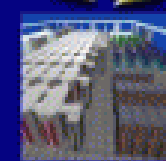
Bio-remediation Carbon Sequestration Energy Production



Distributed Databases



Advanced
Computational
Algorithms

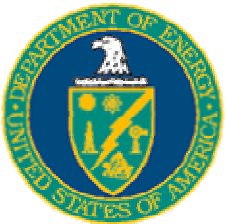


Supercomputer
Resources



Simulation
and Modeling

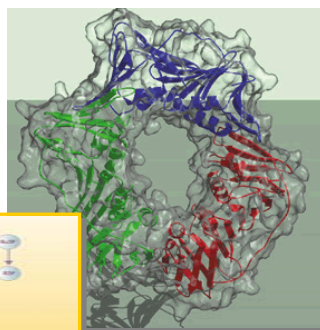
High-performance computing is
essential to high-throughput biology



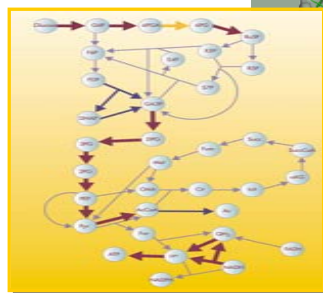
GTL Bridges Physical, Computational and Life sciences

Four research areas:

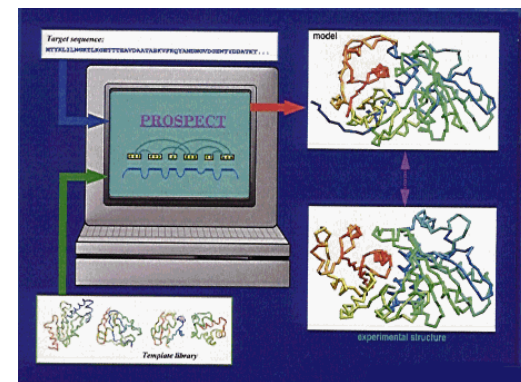
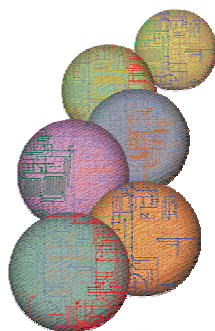
Molecular Machines



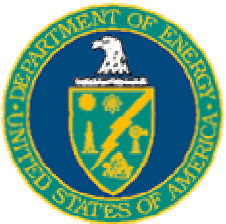
*Regulatory
Pathways*



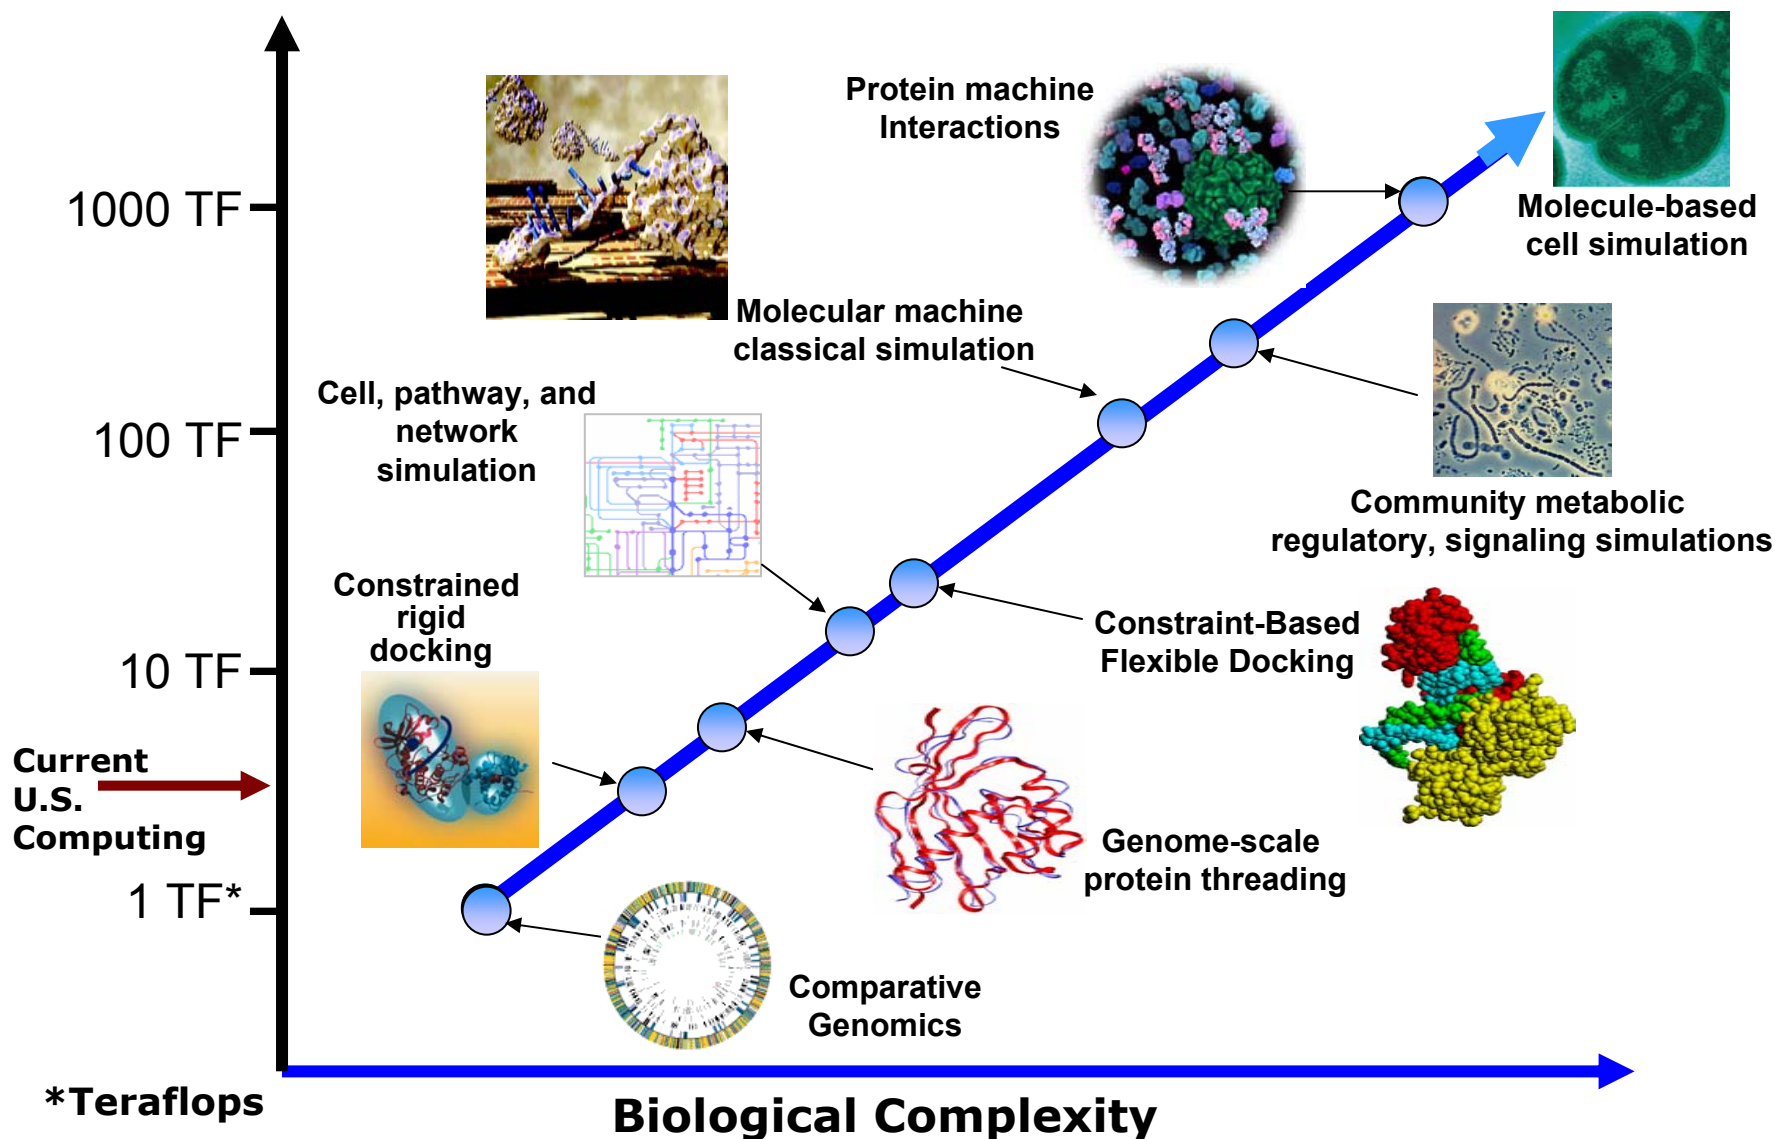
*Microbial
Communities*



*Develop the computational
methods and capabilities to
advance understanding of
complex biological systems*



GTL High-Performance Computing Roadmap



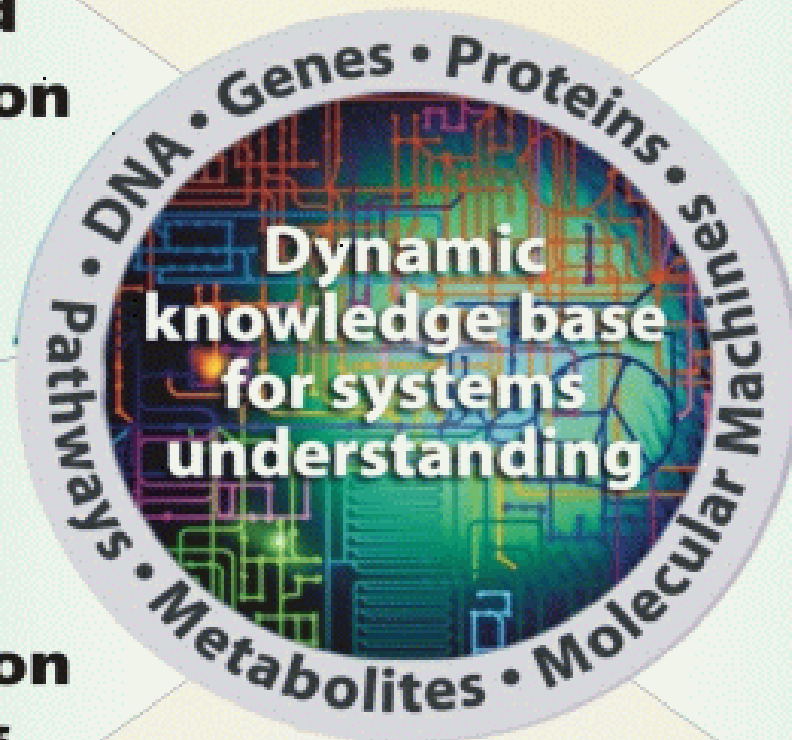
User Facilities for 21st Century Biology

Facility I
Production and
Characterization
of Proteins

Facility II
Whole
Proteome
Analysis

Facility III
Characterization
and Imaging of
Molecular
Machines

Facility IV
Analysis and
Modeling of
Cellular
Systems



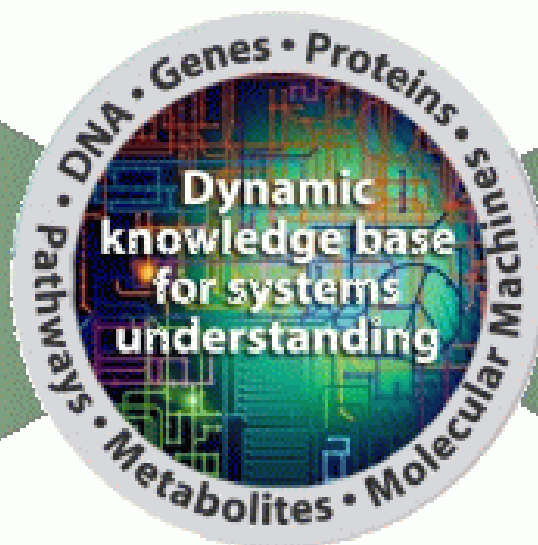


GTL User Facilities Hallmarks

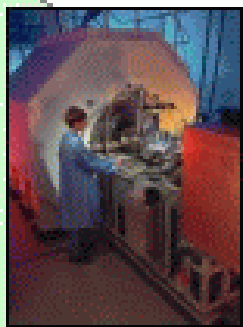


Open, Rapid User Access to Data and Facilities

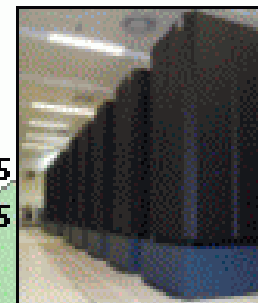
Technology Integration for Tomorrow's Biology



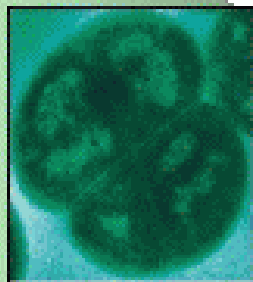
Advanced Technologies



Informatics and Databases



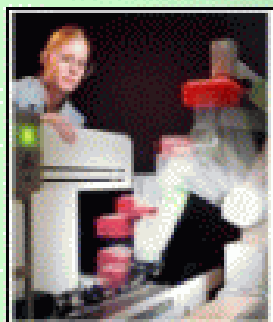
Microbial Genomics



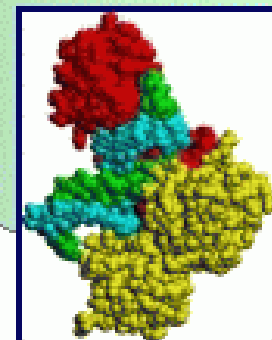
Biosystems Training

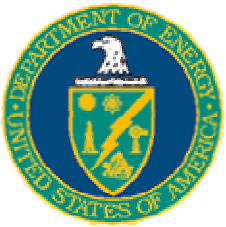


High-Throughput Analysis and Production



Computing and Simulation



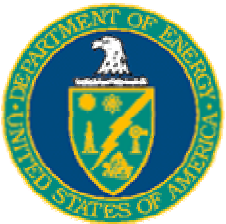


Next Generation Computer Architecture

Goal: Identify and address major hardware and software architectural bottlenecks to the performance of existing and planned DOE science applications

Main Activities

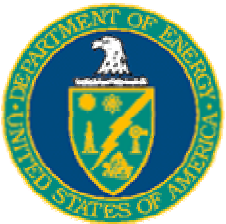
- **Architecture impacts on application performance**
- **OS/runtime research**
- **Evaluation testbeds**



Evaluation of Cray X1 for DOE applications



Delivery to ORNL March 18th
32-processor, liquid-cooled cabinet
128 GB memory
8 TB disk

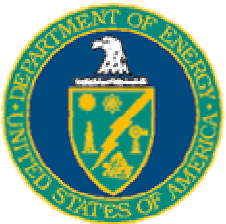


Cray X1 Scale Evaluation


- **Delivery summer 2003**

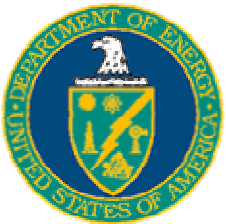


- **3.2 TFlops**
- **256 processors**
- **1 TB shared memory**
- **32 TB of disk space**
- **8 cabinets**

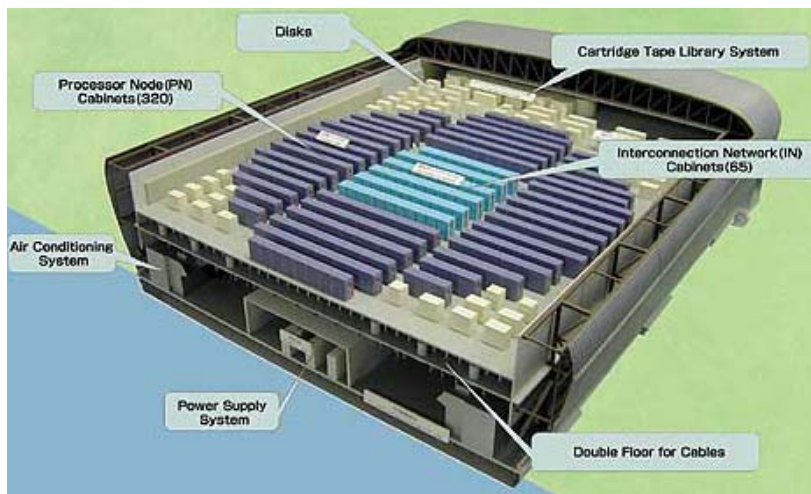


Preliminary X1 evaluation promising

| | SX6 (MFlops) | X1 (MFlops) |
|---|--------------|-------------|
| Climate code  | 3,400 | 5,100 |
| Matrix Multiply | 4,000 | 10,000 |
| Ocean Model | | |
| Kernel A | 3,900 | 4,400 |
| Kernel B | 2,500 | 2,700 |
| Kernel C | 1,000 | 700 |
| Kernel D | 1,800 | 1,400 |
| FE Weather Code | | |
| Kernel A | 84 | 540 |
| Kernel B | 2,000 | 5,900 |
| Kernel C | 150 | 5,300 |
| FE Weather Code | | |
| Kernel A | 2,700 | 4,900 |



16P X1 provides best POP performance on any U.S. Scientific Computer

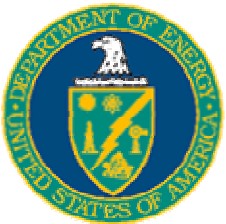


| ES40 | Res | Y/day | Nproc |
|------|-----|-------|-------|
| | 1 | 60.0 | 32 |
| | 0.1 | 3.6 | 960 |

| IBMp4 | Res | Y/day | Nproc |
|-------|-----|-------|-------|
| | 1 | 24.8 | 256 |
| | 0.1 | 0.118 | 480 |

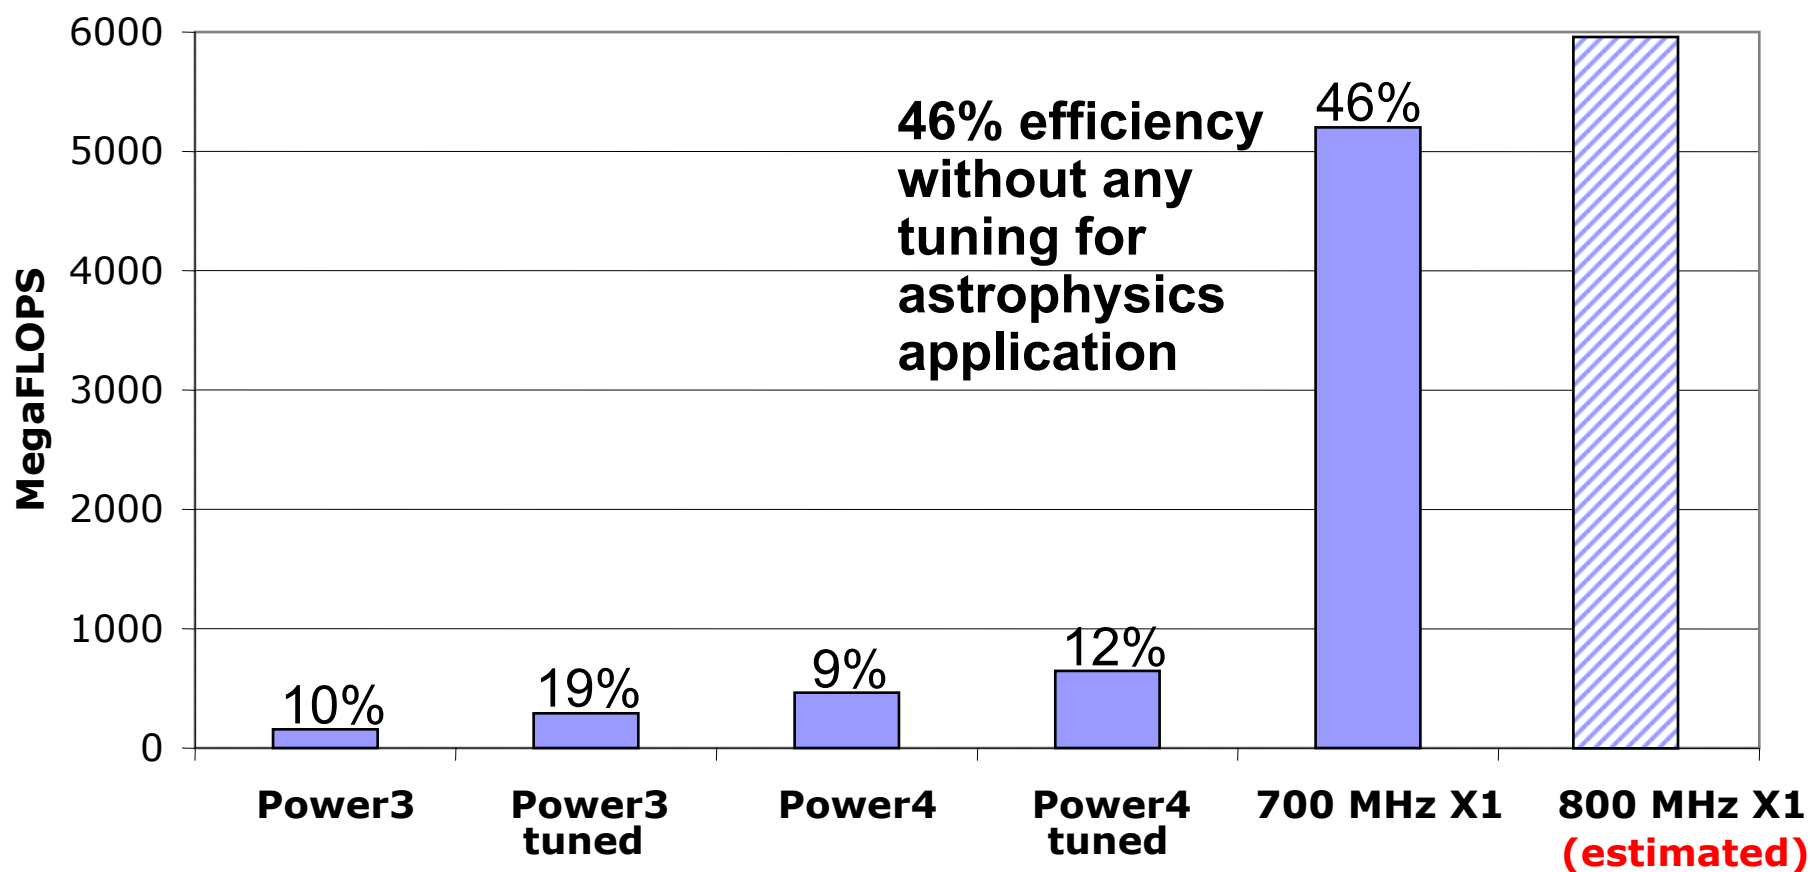


| CrayX1 | Res | Y/day | Nproc |
|--------|-----|-------|-------|
| | 1 | 35.3 | 16 |
| | 0.1 | 0.25 | 16 |

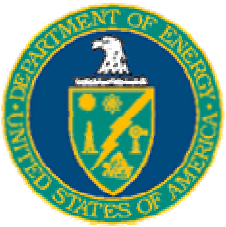


X1 Performance on Boltztran

Boltzman neutrino transport



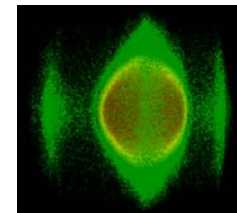
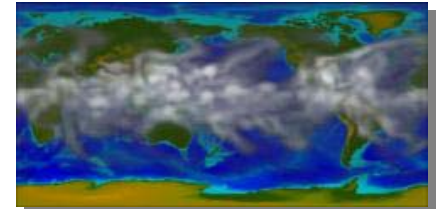
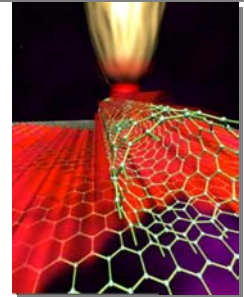
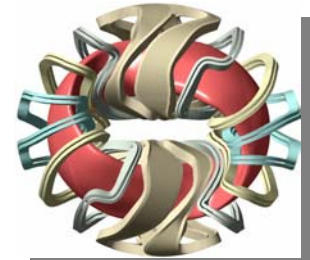
<http://www.phy.ornl.gov/tsi/>

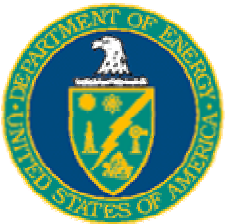


UltraScale Simulation Computing Capability



- **Mission need: Energy production, novel materials, climate science, biological systems**
 - Systems too complex for direct calculation; descriptive laws absent.
 - Involve physical scales up to 50 orders of magnitude;
 - Several scientific disciplines, e.g., combustion; materials science
 - Experimental data may be costly to develop, insufficient, inadequate or unavailable; and
 - Large data files (millions of gigabytes) shared among scientists throughout the world.
- **History of Accomplishments**
 - MPI, Math libraries, first dedicated high-performance computing center, SciDAC





ASCAC Statement

“Without robust response to Earth Simulator, U.S. is open to losing its leadership in defining and advancing frontiers of computational science as new approach to science. This area is critical to both our national security and economic vitality.” (Advanced Scientific Computing Advisory Committee – May 21, 2002).

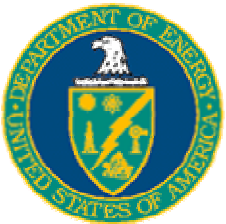
Networking is an important distinguishing characteristic between US needs and Earth Simulator



Simulation Capability Needs

FY2004-05 Timeframe **Need this for Networking**

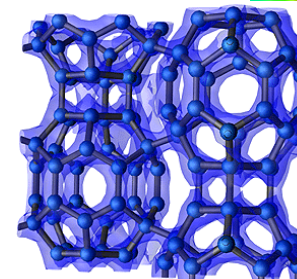
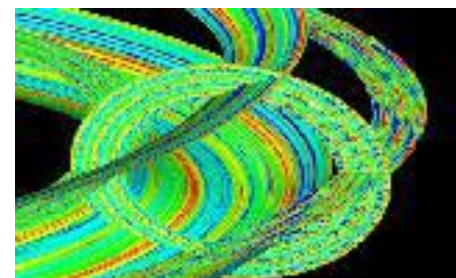
| Application | Simulation Need | Sustained Computational Capability Needed (Tflops) | Significance |
|---------------------------------|--|--|--|
| Climate Science | Calculate chemical balances in atmosphere, including clouds, rivers, and vegetation. | > 50 | Provides U.S. policymakers with leadership data to support policy decisions. Properly represent and predict extreme weather conditions in changing climate. |
| Magnetic Fusion Energy | Optimize balance between self-heating of plasma and heat leakage caused by electromagnetic turbulence. | > 50 | Underpins U.S. decisions about future international fusion collaborations. Integrated simulations of burning plasma crucial for quantifying prospects for commercial fusion. |
| Combustion Science | Understand interactions between combustion and turbulent fluctuations in burning fluid. | > 50 | Understand detonation dynamics (e.g. engine knock) in combustion systems. Solve the "soot" problem in diesel engines. |
| Environmental Molecular Science | Reliably predict chemical and physical properties of radioactive substances. | > 100 | Develop innovative technologies to remediate contaminated soils and groundwater. |
| Astrophysics | Realistically simulate the explosion of a supernova for first time. | >> 100 | Measure size and age of Universe and rate of expansion of Universe. Gain insight into inertial fusion processes. |

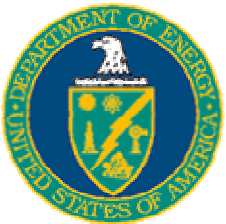


Ultrascalse Scientific Computing

Key Ideas

- Deliver a full-suite of leadership class computers for science with broad applicability.
- Establish a model for computational sciences (SciDAC and base programs) that couples applications scientists, mathematicians, and computational and computer scientists *with computer designers, engineers, and semiconductor researchers.*
- Develop partnerships with domestic computer vendors to ensure that leadership class computers are designed, developed, and produced with science needs as an explicit design criterion.
- Partner with other agencies.
- Partner with industry on applications.





UltraScale Scientific Computing Supporting R&D

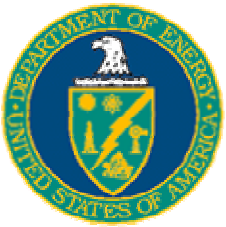


Research with Domestic Vendors – Develop ultrascale hardware and software capabilities for advancing science, focusing on faster interconnects and switches.

Operating Systems, Software Environments, and Tools

- Ensure scalability of operating systems to meet science needs
- Develop enhanced numerical libraries for scientific simulations
- Develop tools to analyze application performance on ultrascale computers

University-based Computer Architecture Research – Explore future generations of computer architectures for ultrascale science simulation.



UltraScale Scientific Computing Computing and Network Facilities

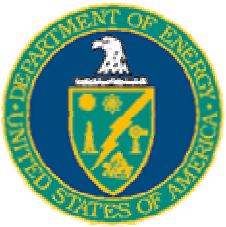


Evaluate computer architectures at levels to ensure that computer hardware and systems software balanced for science and likely to successfully scale

- Continue partnership established between ORNL and Cray, Inc.
- Initiate one new partnership, comprised of scientists and engineers from a domestic computer vendor, with computer scientists, and applications scientists supported by the Office of Science.

Begin installation of first ultrascale computing system for science

Where is the network in this slide?
This is our job – to figure where it fits



To Be Provocative...

Citation in the Press, March 18th, 2008

National Report

The New York Times

DOE Supercomputers Sit Idle

WASHINGTON, Mar. 18, 2008

GAO reports that after almost 5 years of effort and several hundreds of M\$'s spent at the DOE labs, the high performance computers recently purchased did not meet users' expectation and are sitting idle because of lack of network support from ESnet...



How could this happen?

- Users can't get to the resources, cyber security dominates.
- Scientists can't get to the results the supercomputers generated.
- Network unprepared for the sudden increase in computational needs
- Little effort was spent to carry out medium and long term research activities to solve problems that were foreseen 5 years ago, ...